


Hot Topics

LNCS 3405

Alejandro López-Ortiz
Angèle Hamel (Eds.)

Combinatorial and Algorithmic Aspects of Networking

First Workshop on Combinatorial and
Algorithmic Aspects of Networking, CAAN 2004
Banff, Alberta, Canada, August 2004, Revised Selected Papers



Alejandro López-Ortiz Angèle Hamel (Eds.)

Combinatorial and Algorithmic Aspects of Networking

First Workshop on Combinatorial and
Algorithmic Aspects of Networking, CAAN 2004
Banff, Alberta, Canada, August 5-7, 2004
Revised Selected Papers



Volume Editors

Alejandro López-Ortiz
University of Waterloo
School of Computer Science
200 University Ave. W.
Waterloo, Ontario N2L 3G1, Canada
E-mail: alopez-o@uwaterloo.ca

Angèle Hamel
Wilfrid Laurier University
Department of Physics and Computer Science
Waterloo, Ontario N2L 3C5, Canada
E-mail: ahamel@wlu.ca

Library of Congress Control Number: 2005929059

CR Subject Classification (1998): F.1.1, F.2.1-2, C.2, G.2.1-2, E.1

ISSN	0302-9743
ISBN-10	3-540-27873-7 Springer Berlin Heidelberg New York
ISBN-13	978-3-540-27873-3 Springer Berlin Heidelberg New York

This work is subject to copyright. All rights are reserved, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, re-use of illustrations, recitation, broadcasting, reproduction on microfilms or in any other way, and storage in data banks. Duplication of this publication or parts thereof is permitted only under the provisions of the German Copyright Law of September 9, 1965, in its current version, and permission for use must always be obtained from Springer. Violations are liable to prosecution under the German Copyright Law.

Springer is a part of Springer Science+Business Media
springeronline.com

© Springer-Verlag Berlin Heidelberg 2005
Printed in Germany

Typesetting: Camera-ready by author, data conversion by Scientific Publishing Services, Chennai, India
Printed on acid-free paper SPIN: 11527954 06/3142 5 4 3 2 1 0

Commenced Publication in 1973

Founding and Former Series Editors:

Gerhard Goos, Juris Hartmanis, and Jan van Leeuwen

Editorial Board

David Hutchison

Lancaster University, UK

Takeo Kanade

Carnegie Mellon University, Pittsburgh, PA, USA

Josef Kittler

University of Surrey, Guildford, UK

Jon M. Kleinberg

Cornell University, Ithaca, NY, USA

Friedemann Mattern

ETH Zurich, Switzerland

John C. Mitchell

Stanford University, CA, USA

Moni Naor

Weizmann Institute of Science, Rehovot, Israel

Oscar Nierstrasz

University of Bern, Switzerland

C. Pandu Rangan

Indian Institute of Technology, Madras, India

Bernhard Steffen

University of Dortmund, Germany

Madhu Sudan

Massachusetts Institute of Technology, MA, USA

Demetri Terzopoulos

New York University, NY, USA

Doug Tygar

University of California, Berkeley, CA, USA

Moshe Y. Vardi

Rice University, Houston, TX, USA

Gerhard Weikum

Max-Planck Institute of Computer Science, Saarbruecken, Germany

Lecture Notes in Computer Science

For information about Vols. 1–3492

please contact your bookseller or Springer

- Vol. 3626: B. Ganter, G. Stumme, R. Wille (Eds.), *Formal Concept Analysis*. X, 349 pages. 2005. (Subseries LNAI).
- Vol. 3596: F. Dau, M.-L. Mugnier, G. Stumme (Eds.), *Conceptual Structures: Common Semantics for Sharing Knowledge*. XI, 467 pages. 2005. (Subseries LNAI).
- Vol. 3587: P. Perner, A. Imiya (Eds.), *Machine Learning and Data Mining in Pattern Recognition*. XVII, 695 pages. 2005. (Subseries LNAI).
- Vol. 3582: J. Fitzgerald, I.J. Hayes, A. Tarlecki (Eds.), *FM 2005: Formal Methods*. XIV, 558 pages. 2005.
- Vol. 3580: L. Caires, G.F. Italiano, L. Monteiro, C. Palamidessi, M. Yung (Eds.), *Automata, Languages and Programming*. XXV, 1477 pages. 2005.
- Vol. 3578: M. Gallagher, J. Hogan, F. Maire (Eds.), *Intelligent Data Engineering and Automated Learning - IDEAL 2005*. XVI, 599 pages. 2005.
- Vol. 3576: K. Etessami, S.K. Rajamani (Eds.), *Computer Aided Verification*. XV, 564 pages. 2005.
- Vol. 3575: S. Wermter, G. Palm, M. Elshaw (Eds.), *Biomimetic Neural Learning for Intelligent Robots*. IX, 383 pages. 2005. (Subseries LNAI).
- Vol. 3574: C. Boyd, J.M. González Nieto (Eds.), *Information Security and Privacy*. XIII, 586 pages. 2005.
- Vol. 3573: S. Etalle (Ed.), *Logic Based Program Synthesis and Transformation*. VIII, 279 pages. 2005.
- Vol. 3572: C. De Felice, A. Restivo (Eds.), *Developments in Language Theory*. XI, 409 pages. 2005.
- Vol. 3571: L. Godo (Ed.), *Symbolic and Quantitative Approaches to Reasoning with Uncertainty*. XVI, 1028 pages. 2005. (Subseries LNAI).
- Vol. 3570: A. S. Patrick, M. Yung (Eds.), *Financial Cryptography and Data Security*. XII, 376 pages. 2005.
- Vol. 3569: F. Bacchus, T. Walsh (Eds.), *Theory and Applications of Satisfiability Testing*. XII, 492 pages. 2005.
- Vol. 3568: W.K. Leow, M.S. Lew, T.-S. Chua, W.-Y. Ma, L. Chaisorn, E.M. Bakker (Eds.), *Image and Video Retrieval*. XVII, 672 pages. 2005.
- Vol. 3567: M. Jackson, D. Nelson, S. Stirk (Eds.), *Database: Enterprise, Skills and Innovation*. XII, 185 pages. 2005.
- Vol. 3566: J.-P. Banâtre, P. Fradet, J.-L. Giavitto, O. Michel (Eds.), *Unconventional Programming Paradigms*. XI, 367 pages. 2005.
- Vol. 3565: G.E. Christensen, M. Sonka (Eds.), *Information Processing in Medical Imaging*. XXI, 777 pages. 2005.
- Vol. 3564: N. Eisinger, J. Małuszyński (Eds.), *Reasoning Web*. IX, 319 pages. 2005.
- Vol. 3562: J. Mira, J.R. Álvarez (Eds.), *Artificial Intelligence and Knowledge Engineering Applications: A Bioinspired Approach*, Part II. XXIV, 636 pages. 2005.
- Vol. 3561: J. Mira, J.R. Álvarez (Eds.), *Mechanisms, Symbols, and Models Underlying Cognition*, Part I. XXIV, 532 pages. 2005.
- Vol. 3560: V.K. Prasanna, S. Iyengar, P.G. Spirakis, M. Welsh (Eds.), *Distributed Computing in Sensor Systems*. XV, 423 pages. 2005.
- Vol. 3559: P. Auer, R. Meir (Eds.), *Learning Theory*. XI, 692 pages. 2005. (Subseries LNAI).
- Vol. 3558: V. Torra, Y. Narukawa, S. Miyamoto (Eds.), *Modeling Decisions for Artificial Intelligence*. XII, 470 pages. 2005. (Subseries LNAI).
- Vol. 3557: H. Gilbert, H. Handschuh (Eds.), *Fast Software Encryption*. XI, 443 pages. 2005.
- Vol. 3556: H. Baumeister, M. Marchesi, M. Holcombe (Eds.), *Extreme Programming and Agile Processes in Software Engineering*. XIV, 332 pages. 2005.
- Vol. 3555: T. Vardanega, A. Wellings (Eds.), *Reliable Software Technology – Ada-Europe 2005*. XV, 273 pages. 2005.
- Vol. 3554: A. Dey, B. Kokinov, D. Leake, R. Turner (Eds.), *Modeling and Using Context*. XIV, 572 pages. 2005. (Subseries LNAI).
- Vol. 3553: T.D. Hämmäläinen, A.D. Pimentel, J. Takala, S. Vassiliadis (Eds.), *Embedded Computer Systems: Architectures, Modeling, and Simulation*. XV, 476 pages. 2005.
- Vol. 3552: H. de Meer, N. Bhatti (Eds.), *Quality of Service – IWQoS 2005*. XVIII, 400 pages. 2005.
- Vol. 3551: T. Härder, W. Lehner (Eds.), *Data Management in a Connected World*. XIX, 371 pages. 2005.
- Vol. 3548: K. Julisch, C. Kruegel (Eds.), *Intrusion and Malware Detection and Vulnerability Assessment*. X, 241 pages. 2005.
- Vol. 3547: F. Bomarius, S. Komi-Sirviö (Eds.), *Product Focused Software Process Improvement*. XIII, 588 pages. 2005.
- Vol. 3546: T. Kanade, A. Jain, N.K. Ratha (Eds.), *Audio- and Video-Based Biometric Person Authentication*. XX, 1134 pages. 2005.
- Vol. 3543: L. Kutvonen, N. Alonistioti (Eds.), *Distributed Applications and Interoperable Systems*. XI, 235 pages. 2005.
- Vol. 3542: H.H. Hoos, D.G. Mitchell (Eds.), *Theory and Applications of Satisfiability Testing*. XIII, 393 pages. 2005.
- Vol. 3541: N.C. Oza, R. Polikar, J. Kittler, F. Roli (Eds.), *Multiple Classifier Systems*. XII, 430 pages. 2005.

- Vol. 3540: H. Kalviainen, J. Parkkinen, A. Kaarna (Eds.), *Image Analysis*. XXII, 1270 pages. 2005.
- Vol. 3537: A. Apostolico, M. Crochemore, K. Park (Eds.), *Combinatorial Pattern Matching*. XI, 444 pages. 2005.
- Vol. 3536: G. Ciardo, P. Darondeau (Eds.), *Applications and Theory of Petri Nets* 2005. XI, 470 pages. 2005.
- Vol. 3535: M. Steffen, G. Zavattaro (Eds.), *Formal Methods for Open Object-Based Distributed Systems*. X, 323 pages. 2005.
- Vol. 3533: M. Ali, F. Esposito (Eds.), *Innovations in Applied Artificial Intelligence*. XX, 858 pages. 2005. (Subseries LNAI).
- Vol. 3532: A. Gómez-Pérez, J. Euzenat (Eds.), *The Semantic Web: Research and Applications*. XV, 728 pages. 2005.
- Vol. 3531: J. Ioannidis, A. Keromytis, M. Yung (Eds.), *Applied Cryptography and Network Security*. XI, 530 pages. 2005.
- Vol. 3530: A. Prinz, R. Reed, J. Reed (Eds.), *SDL 2005: Model Driven*. XI, 361 pages. 2005.
- Vol. 3528: P.S. Szczepaniak, J. Kacprzyk, A. Niewiadomski (Eds.), *Advances in Web Intelligence*. XVII, 513 pages. 2005. (Subseries LNAI).
- Vol. 3527: R. Morrison, F. Oquendo (Eds.), *Software Architecture*. XII, 263 pages. 2005.
- Vol. 3526: S.B. Cooper, B. Löwe, L. Torenvliet (Eds.), *New Computational Paradigms*. XVII, 574 pages. 2005.
- Vol. 3525: A.E. Abdallah, C.B. Jones, J.W. Sanders (Eds.), *Communicating Sequential Processes*. XIV, 321 pages. 2005.
- Vol. 3524: R. Barták, M. Milano (Eds.), *Integration of AI and OR Techniques in Constraint Programming for Combinatorial Optimization Problems*. XI, 320 pages. 2005.
- Vol. 3523: J.S. Marques, N. Pérez de la Blanca, P. Pina (Eds.), *Pattern Recognition and Image Analysis, Part II*. XXVI, 733 pages. 2005.
- Vol. 3522: J.S. Marques, N. Pérez de la Blanca, P. Pina (Eds.), *Pattern Recognition and Image Analysis, Part I*. XXVI, 703 pages. 2005.
- Vol. 3521: N. Megiddo, Y. Xu, B. Zhu (Eds.), *Algorithmic Applications in Management*. XIII, 484 pages. 2005.
- Vol. 3520: O. Pastor, J. Falcão e Cunha (Eds.), *Advanced Information Systems Engineering*. XVI, 584 pages. 2005.
- Vol. 3519: H. Li, P. J. Olver, G. Sommer (Eds.), *Computer Algebra and Geometric Algebra with Applications*. IX, 449 pages. 2005.
- Vol. 3518: T.B. Ho, D. Cheung, H. Liu (Eds.), *Advances in Knowledge Discovery and Data Mining*. XXI, 864 pages. 2005. (Subseries LNAI).
- Vol. 3517: H.S. Baird, D.P. Lopresti (Eds.), *Human Interactive Proofs*. IX, 143 pages. 2005.
- Vol. 3516: V.S. Sunderam, G.D.v. Albada, P.M.A. Sloot, J.J. Dongarra (Eds.), *Computational Science – ICCS 2005, Part III*. LXIII, 1143 pages. 2005.
- Vol. 3515: V.S. Sunderam, G.D.v. Albada, P.M.A. Sloot, J.J. Dongarra (Eds.), *Computational Science – ICCS 2005, Part II*. LXIII, 1101 pages. 2005.
- Vol. 3514: V.S. Sunderam, G.D.v. Albada, P.M.A. Sloot, J.J. Dongarra (Eds.), *Computational Science – ICCS 2005, Part I*. LXIII, 1089 pages. 2005.
- Vol. 3513: A. Montoyo, R. Muñoz, E. Métais (Eds.), *Natural Language Processing and Information Systems*. XII, 408 pages. 2005.
- Vol. 3512: J. Cabestany, A. Prieto, F. Sandoval (Eds.), *Computational Intelligence and Bioinspired Systems*. XXV, 1260 pages. 2005.
- Vol. 3511: U.K. Wiil (Ed.), *Metainformatics*. VIII, 221 pages. 2005.
- Vol. 3510: T. Braun, G. Carle, Y. Koucheryavy, V. Tsoulos (Eds.), *Wired/Wireless Internet Communications*. XIV, 366 pages. 2005.
- Vol. 3509: M. Jünger, V. Kaibel (Eds.), *Integer Programming and Combinatorial Optimization*. XI, 484 pages. 2005.
- Vol. 3508: P. Bresciani, P. Giorgini, B. Henderson-Sellers, G. Low, M. Winikoff (Eds.), *Agent-Oriented Information Systems II*. X, 227 pages. 2005. (Subseries LNAI).
- Vol. 3507: F. Crestani, I. Ruthven (Eds.), *Information Context: Nature, Impact, and Role*. XIII, 253 pages. 2005.
- Vol. 3506: C. Park, S. Chee (Eds.), *Information Security and Cryptology – ICISC 2004*. XIV, 490 pages. 2005.
- Vol. 3505: V. Gorodetsky, J. Liu, V. A. Skormin (Eds.), *Autonomous Intelligent Systems: Agents and Data Mining*. XIII, 303 pages. 2005. (Subseries LNAI).
- Vol. 3504: A.F. Frangi, P.I. Radeva, A. Santos, M. Hernandez (Eds.), *Functional Imaging and Modeling of the Heart*. XV, 489 pages. 2005.
- Vol. 3503: S.E. Nikolettas (Ed.), *Experimental and Efficient Algorithms*. XV, 624 pages. 2005.
- Vol. 3502: F. Khendek, R. Dssouli (Eds.), *Testing of Communicating Systems*. X, 381 pages. 2005.
- Vol. 3501: B. Kégl, G. Lapalme (Eds.), *Advances in Artificial Intelligence*. XV, 458 pages. 2005. (Subseries LNAI).
- Vol. 3500: S. Miyano, J. Mesirov, S. Kasif, S. Istrail, P. Pevzner, M. Waterman (Eds.), *Research in Computational Molecular Biology*. XVII, 632 pages. 2005. (Subseries LNBI).
- Vol. 3499: A. Pelc, M. Raynal (Eds.), *Structural Information and Communication Complexity*. X, 323 pages. 2005.
- Vol. 3498: J. Wang, X. Liao, Z. Yi (Eds.), *Advances in Neural Networks – ISNN 2005, Part III*. XLIX, 1077 pages. 2005.
- Vol. 3497: J. Wang, X. Liao, Z. Yi (Eds.), *Advances in Neural Networks – ISNN 2005, Part II*. XLIX, 947 pages. 2005.
- Vol. 3496: J. Wang, X. Liao, Z. Yi (Eds.), *Advances in Neural Networks – ISNN 2005, Part I*. L, 1055 pages. 2005.
- Vol. 3495: P. Kantor, G. Muresan, F. Roberts, D.D. Zeng, F.-Y. Wang, H. Chen, R.C. Merkle (Eds.), *Intelligence and Security Informatics*. XVIII, 674 pages. 2005.
- Vol. 3494: R. Cramer (Ed.), *Advances in Cryptology – EUROCRYPT 2005*. XIV, 576 pages. 2005.
- Vol. 3493: N. Fuhr, M. Lalmas, S. Malik, Z. Szilávik (Eds.), *Advances in XML Information Retrieval*. XI, 438 pages. 2005.

Preface

The Internet is a massive global network of over 700 million users and it is adding users at the rate of 300,000 per day. This large, distributed, and everchanging network poses a challenge to researchers: How does one study, model, or understand such a decentralized, constantly evolving entity? Research in large-scale networks seeks to address this question, and the unique nature of these networks calls for a range of techniques from a host of disciplines. The workshop Combinatorial and Algorithmic Aspects of Networking and the Internet (CAAN 2004) provided a forum for the exchange of ideas on these topics.

The primary goals of the workshop were to bring together a diverse cross-section of researchers in an already scattered and distinct community and also to provide a snapshot of the cutting-edge research in this field. We succeeded in these goals: among the participants were mathematicians, computer scientists in theory and algorithms, computer scientists in networks, physicists, and engineers, as well as researchers from Europe and North America, participants from industry and academia, students, and established researchers; and among the papers were some new and surprising results as well as some introductions to the foundations of the field.

The workshop program featured 12 peer-reviewed papers bracketed by two hour-long invited survey talks—an opening talk by Ashish Goel and a closing talk by Andrei Broder. Topics covered by the talks ranged from the Web graph to game theory to string matching, all in the context of large-scale networks. This volume collects together the talks delivered at the workshop along with a number of survey articles to round out the presentation and give a comprehensive introduction to the topic.

We were fortunate to be given the opportunity to hold the conference as a two-day workshop at the Banff International Research Station for Mathematical Innovation and Discovery, BIRS. The breathtaking and inspiring setting and ample amenities contributed greatly to the success of the workshop. Attendance at BIRS is by invitation only and we had 24 participants at CAAN 2004. The small number of participants facilitated an intimate atmosphere perfect for generating discussions and initiating collaborations.

We would like to thank the Steering Committee for their guidance, and the Program Committee for their diligent work in reviewing the papers and selecting an excellent and balanced program. Special thanks goes to the organizational team at BIRS, especially Andrea Lundquist and Jackie Kler who made our job easy. We also thank Chris Taylor (<http://photos.t-a-y-l-o-r.com>) who graciously gave us permission to use the photographic artwork on the workshop poster. Finally, of course, we thank all the participants whose enthusiasm and support made CAAN 2004 a success and encouraged us to offer the workshop again.

CAAN 2005 will be held in August 2005 in Waterloo, Ontario, Canada as a satellite workshop of the Workshop on Algorithms and Data Structures (WADS 2005).

Organization

Steering Committee

Andrei Broder
Angèle Hamel
Srinivasan Keshav
Alejandro López-Ortiz
Rajeev Motwani
Ian Munro

IBM, USA
Wilfrid Laurier University, Canada
University of Waterloo, Canada
University of Waterloo, Canada
Stanford University, USA
University of Waterloo, Canada

Program Committee

Azer Bestavros
Anthony Bonato
Andrei Broder
Hervé Brönnimann
Adam Buchsbaum
Edith Cohen
Erik Demaine
Luisa Gargano
Ashish Goel
Angèle Hamel (Org. Comm. Chair)
Monika Henzinger
Jeannette Janssen
David Karger
Srinivasan Keshav
Alejandro López-Ortiz (Program Chair)
Bruce Maggs
Rajeev Motwani
Ian Munro
Tim Roughgarden

Christos Papadimitrou

David Peleg
Torsten Suel
Eli Upfal
Alessandro Vespignani

Boston University, USA
Wilfrid Laurier University, Canada
IBM Research, USA
Polytechnic University, USA
AT&T Labs, USA
AT&T Labs, USA
MIT, USA
University of Salerno, Italy
Stanford University, USA
Wilfrid Laurier University, Canada
Google, USA
Dalhousie University, Canada
MIT, USA
University of Waterloo, Canada
University of Waterloo, Canada
Carnegie Mellon University, USA
Stanford University, USA
University of Waterloo, Canada
University of California at Berkeley,
USA
University of California at Berkeley,
USA
Weizmann Institute, Israel
Polytechnic University, USA
Brown University, USA
Laboratoire de Physique
Théorique, France

Table of Contents

Workshop Proceedings

Aggregating Correlated Data in Sensor Networks <i>Mihaela Enachescu, Ashish Goel, Ramesh Govindan, Rajeev Motwani</i>	1
The Efficiency of Optimal Taxes <i>George Karakostas, Stavros G. Kolliopoulos</i>	3
Congestion Games, Load Balancing, and Price of Anarchy <i>Anshul Kothari, Subhash Suri, Csaba D. Tóth, Yunhong Zhou</i>	13
Bandwidth Allocation in Networks: A Single Dual Update Subroutine for Multiple Objectives <i>Sung-woo Cho, Ashish Goel</i>	28
Limits and Power Laws of Models for the Web Graph and Other Networked Information Spaces <i>Anthony Bonato, Jeannette Janssen</i>	42
Cuts and Disjoint Paths in the Valley-Free Path Model of Internet BGP Routing <i>Thomas Erlebach, Alexander Hall, Alessandro Panconesi, Danica Vukadinović</i>	49
A Distributed Algorithm to Find Hamiltonian Cycles in $\mathcal{G}(n, p)$ Random Graphs <i>Eythan Levy, Guy Louchard, Jordi Petit</i>	63
String Matching on the Internet <i>Hervé Brönnimann, Nasir Memon, Kulesh Shanmugasundaram</i>	75
k -Robust Single-Message Transmission <i>André Kündgen, Michael J. Pelsmayer, Radhika Ramamurthi</i>	90
Stable Local Scheduling Algorithms With Low Complexity and Without Speedup for a Network of Input-Buffered Switches <i>Claus Bauer</i>	102
The External Network Problem with Edge- or Arc-Connectivity Requirements <i>Jan van den Heuvel, Matthew Johnson</i>	114

VIII Table of Contents

Bipartite Graphs as Models of Complex Networks <i>Jean-Loup Guillaume, Matthieu Latapy</i>	127
Traceroute-Like Exploration of Unknown Networks: A Statistical Analysis <i>Luca Dall'Asta, Ignacio Alvarez-Hamelin, Alain Barrat, Alexei Vázquez, Alessandro Vespignani</i>	140
Invited Talk: The Many Wonders of the Web Graph <i>Andrei Broder</i>	154
Survey Articles	
Algorithmic Foundations of the Internet: Foreword <i>Alejandro López-Ortiz</i>	155
A Survey of Models of the Web Graph <i>Anthony Bonato</i>	159
You Can Get There from Here: Routing in the Internet <i>Jean-Charles Grégoire, Angèle M. Hamel</i>	173
Search Engines and Web Information Retrieval <i>Alejandro López-Ortiz</i>	183
Algorithmic Foundations of the Internet: Roundup <i>Alejandro López-Ortiz</i>	192
Author Index	205

Aggregating Correlated Data in Sensor Networks

M. Enachescu¹, A. Goel¹, R. Govindan², and R. Motwani¹

¹ Stanford University, California, USA

² University of Southern California, California, USA

Abstract for Invited Presentation. Consider a network where each node gathers information from its vicinity and sends this information to a centralized processing agent. If the information is geographically correlated, then a large saving in data transmission costs may be obtained by aggregating information from nearby nodes before sending it to the central agent. This is particularly relevant to sensor networks where battery limitations dictate that data transmission be kept to a minimum, and where sensed data is often geographically correlated. In-network aggregation for sensor networks has been extensively studied over the last few years. In this paper we show that a very simple opportunistic aggregation scheme can result in near-optimum performance under widely varying (and unknown) scales of correlation.

More formally, we consider the idealized setting where sensors are arranged on an $N \times N$ grid, and the centralized processing agent is located at position $(0,0)$ on the grid. We assume that each sensor can communicate only to its four neighbors on the grid. This idealized setting has been widely used to study broad information processing issues in sensor networks. We call an aggregation scheme *opportunistic* if data from a sensor to the central agent is always sent over a shortest path, i.e., no extra routing penalty is incurred to achieve aggregation.

To model geographic correlations, we assume that each sensor can gather information in a $\frac{k}{2} \times \frac{k}{2}$ square (or, a circle of radius $k/2$) centered at the sensor. We will refer to k as the correlation parameter. Let the set $A(x)$ denote the area sensed by sensor x . If we aggregate information from a set of sensors S then the size of the resulting compressed information is $I(S) = |\bigcup_{x \in S} A(x)|$, i.e., the size of the total area covered by the sensors in S . Often, the parameter k can depend on the intensity of the information being sensed. For example, a volcanic eruption might be recorded by many more sensors and would correspond to a much higher k than a campfire. Accordingly, we will assume that the parameter k is not known in advance. In fact, we would like our opportunistic aggregation algorithms to work well simultaneously for all k .

There are scenarios where the above model applies directly. For example, the sensors could be cameras which take pictures within a certain radius, or they could be sensing RFID tags on retail items (or on birds which have been tagged for environmental monitoring) within a certain radius. Also, since we want algorithms that work well without any knowledge of k , our model applies to scenarios where the likelihood of sensing decreases with distance. Thus, we believe that our model

(optimizing simultaneously for all k) captures the joint entropy of correlated sets of sensors in a natural way for a large variety of applications.

For node (i, j) , we will refer to nodes $(i-1, j)$ and $(i, j-1)$ as its downstream neighbors, and nodes $(i+1, j)$ and $(i, j+1)$ as its upstream neighbors. We would like to construct a tree over which information flows to the central agent, and gets aggregated along the way. Since we are restricted to routing over shortest paths, each node has just one choice: which downstream node to choose as its parent in the tree. Each node aggregates the information it sensed locally with any information it received from its upstream neighbors and sends it on to one of its downstream neighbors. The cost of the tree is the total amount of (compressed) information sent out over links in the tree.

Our Results: We propose a very simple randomized algorithm for choosing the next neighbor – node (i, j) chooses its left neighbor with probability $i/(i+j)$ and its bottom neighbor with probability $j/(i+j)$. Observe that this scheme results in all shortest paths between (i, j) and $(0, 0)$ being chosen with equal probability. We prove that this simple scheme is a constant factor approximation (in expectation) to the optimum aggregation tree *simultaneously* for all correlation parameters k . While we construct a single tree, the optimum trees for different correlation parameters may be different. The key idea is to relate the expected collision time of random walks on the grid to scale free aggregation.

Our results hold only for the total cost, and critically rely on the fact that information is distributed evenly through the sensor field. This result shows that, at least for the class of aggregation functions and the grid topology considered in this paper, schemes that attempt to construct specialized routing structures in order to improve the likelihood of data aggregation are unnecessary. This is convenient, since such specialized routing structures are hard to build without some a priori knowledge about correlations in the data. With this result, simple geographic routing schemes or tree-based data gathering protocols are sufficient.

The Efficiency of Optimal Taxes

George Karakostas^{1,*} and Stavros G. Kolliopoulos^{2,**}

¹ Department of Computing and Software, McMaster University
karakos@mcmaster.ca

² Department of Informatics and Telecommunications,
University of Athens and Department of Computing and Software,
McMaster University
www.cas.mcmaster.ca/~stavros

Abstract. It is well known that the selfish behavior of users in a network can be regulated through the imposition of the so-called *optimal taxes* on the network edges. Any traffic equilibrium reached by the selfish users who are conscious of both the travel latencies and the taxes will minimize the social cost, i.e., will minimize the total latency.

Optimal taxes incur desirable behavior from the society point of view but they cause disutility to the network users since the users' total cost is in general increased [4]. Excessive disutility due to taxation may be undesirable from the societal perspective as well. In this work we examine the efficiency of taxation as a mechanism for achieving the desired goal of minimizing the social cost. We show that for large classes of latency functions the total disutility due to taxation that is caused to the users and/or the system is bounded with respect to the social optimum. In addition, we show that if the social cost takes into account both the total latency and the total taxation in the network, the coordination ratio for certain latency functions is better than the coordination ratio when taxation is not used.

1 Introduction

In the *selfish routing* setting, we are given a directed network $G = (V, E)$ and a set of k classes of users (commodities), each with its own origin and destination, and with a fixed total demand (traffic) rate per class $d_i > 0$, $i = 1, \dots, k$. Individual users are thought as carrying each an infinitesimal amount of a commodity. We are also given a nonnegative latency function l_P describing the delay experienced by users wishing to travel on the path P as a function of the total flow through the edges of the path. In this work we assume that the *additive model* holds, i.e., for every edge e there is a latency function $l_e(f_e)$ that describes the latency on this edge due to the flow f_e that crosses it. Then the latency for a path is defined as $l_P(f) := \sum_{e \in P} l_e(f_e)$. Each user tries to selfishly route his flow so

* Research supported by an NSERC Discovery grant.

** Research partially supported by NSERC Research Grant 227809-00.

that his path cost is minimized. A *traffic equilibrium* is an assignment of traffic to paths so that no user can unilaterally switch her flow to a path of smaller cost. Wardrop's principle [16] for selfish routing postulates that

at equilibrium, for each origin-destination pair, the travel costs on all the routes actually used are equal, or less than the travel costs on all nonused routes.

A widely used measure of the network performance is the *social cost* (or *total latency*), defined as $\sum_{\text{path } P} f_P l_P(f) = \sum_{e \in E} f_e l_e(f_e)$ for a flow f that routes f_P units of traffic through path P . Although it must obey Wardrop's principle at equilibrium, the unregulated choice of paths by individual users may incur a social cost which in general can be higher than the social optimum. In fact the latency of an equilibrium can be arbitrarily larger than the social optimum [14]. A classic way of dealing with this problem is the introduction of *taxation* on the edges of the network, so that the users' path cost has both a travel time and a budgetary component. Without taxation, users experience only their own traffic delay as their cost. With taxation users are also charged for the right to use a path. This technique has been studied by the traffic community for a long time (cf. [5] and the references therein), especially in the context of *marginal costs* (see, for example, [2],[8],[15]). Each selfish user of class i using path P will experience the following path cost:

$$\text{path cost}(P) := \text{latency}(P) + a(i) \cdot \text{taxation}(P).$$

The $\text{taxation}(P)$ is the sum of taxes along the edges of the path. The factor $a(i) > 0$, denotes the sensitivity of user class i to the taxes. In the *homogeneous* case all user classes have the same sensitivity to the taxation (i.e. $a(i) = 1$, for all i), while in the *heterogeneous* case $a(i)$ can take different positive values for different classes. Through edge taxation, we would like to force *all* equilibria on the network to induce flow that minimizes the social cost $\sum_{e \in E} f_e l_e(f_e)$. We refer to a set of edge taxes that achieves this as *optimal taxes*. In the homogeneous case, marginal costs have been shown (cf. [2],[8],[15]) to be optimal taxes. In the heterogeneous case, the existence and calculation of such optimal taxes were shown for the single source-destination pair case by Cole et al. [5], and were later extended to the multicommodity setting in [9, 11].

Designing optimal taxes is a classic instance of mechanism design, a central topic in game theory (see, e.g., [13]). A set of outcomes is fixed (here achieving the social optimum) and users are paid or penalized, (here they pay taxes) in order to achieve the desired outcome in equilibrium. One can actually see the taxation cost in two different contexts. One is the context already discussed, which is as monetary cost. The other is to see the tax for every edge as part of the edge latency function itself. Then, instead of taxation, we can speak about *artificial delays* introduced possibly at the entrance of each edge, in order to minimize the total amount of time users actually spend on the edges themselves. For example, this is the technique used at some highway exits, where traffic lights have been installed in order to better control traffic. Whatever the context

though, taxation increases in general the user cost, as was shown in [4] for the case of marginal cost taxes. The natural question that arises is whether taxes are an efficient mechanism for achieving the desired result. Is the additional disutility caused through taxation proportionate to the desired goal, i.e., a routing that minimizes the total latency? In this paper we tackle this problem by comparing the social cost of the traffic equilibria when taxation is used against (i) the social optimum without taxation and (ii) the social optimum when taxation is taken into account.

In Section 4 we show that in the homogeneous case the ratio of social cost at equilibrium with taxation to the social optimum without taxation is not much bigger than the worst case ratio without any taxation for many important families of latency functions, like linear or low-degree polynomial ones. In particular for strictly increasing linear latency functions we show that, if b is the vector of optimal edge taxes (in this case, the marginal cost taxes),

$$\frac{\sum_e f_e^*(l_e(f_e^*) + b_e)}{\sum_e \hat{f}_e l_e(\hat{f}_e)} \leq 2$$

for any equilibrium flow f^* and flow \hat{f} that achieves the social optimum. This bound is tight, and is not far from the $4/3$ upper bound on the *coordination ratio* in the case without taxes, shown by Roughgarden and Tardos [14]. The coordination ratio ρ was defined by Koutsoupias and Papadimitriou in [12] as follows

$$\rho := \sup_{f^*} \frac{\sum_e f_e^* l_e(f_e^*)}{\sum_e \hat{f}_e l_e(\hat{f}_e)}$$

for the worst case (in terms of social cost) equilibrium f^* , and \hat{f} as before. Hence we show that with a small increase in network inefficiency, we achieve, at equilibrium, a flow pattern that minimizes the total latency of the users. Note that, in principle, the tax b_e on an edge e could be very big compared to the latency part $l_e()$ of the edge cost function. Hence it is rather surprising that taxation does not drive the social cost further than a small constant factor away from the social (without taxation) optimum.

This approach in bounding the inefficiency of taxation as a mechanism to achieve minimum social cost is influenced by the notion of coordination mechanisms. This concept was recently introduced by Christodoulou, Koutsoupias and Nanavati [3]. Informally speaking a coordination mechanism is a cost function experienced by the users, chosen from a family of possible cost functions called a coordination model. The measure of the efficiency of a coordination mechanism is the supremum over all possible demand sets of the ratio of the social cost of the worst-case equilibrium to the social optimum achieved with the *original* cost function. See [3] for the precise mathematical definitions. Note that in our work the demands are fixed.

Proving that taxation incurs a small increase to the cost of an equilibrium compared to the social optimum without taxation is a satisfying result. However it is possible that once taxes are imposed by some central authority they are

considered to be part of the social cost. In other words taxation may incur disutility to society as a whole. To address this issue, we compare the worst-case cost with taxation of an equilibrium against the social optimum with taxation. In other words we consider the coordination ration in the standard sense [12] of the game with taxes. In Section 3 we show that, for certain families of strictly increasing and continuous latency functions (like linear or polynomial ones), the coordination ratio of the network actually *decreases* when optimal taxes are introduced. In particular for the linear latency functions case, we show that

$$\frac{\sum_e f_e^*(l_e(f_e^*) + b_e)}{\sum_e \bar{f}_e(l_e(\bar{f}_e) + b_e)} \leq \frac{5}{4}$$

for any equilibrium flow f^* and social optimum \bar{f} when taxes b are used in both cases. This is significantly better than the $4/3$ bound of [14] for general linear functions. The gap between $4/3$ and $5/4$ quantifies the beneficial effect of taxation on the behavior of the selfish users, specifically the reduction in their resistance to coordination. Our bound holds for heterogeneous users as well, and its proof is based on the definition of the parameter $\beta(\mathcal{L})$ for a family of functions \mathcal{L} by Correa, Schulz and Stier Moses [6].

The two results in combination show that imposing the optimal taxes on a selfish routing game not only yields a game with better coordination ratio, but also the added disutility to the users is bounded with respect to the original system optimum. In addition we emphasize that our two approaches together provide a stronger guarantee on the worst-case cost of an equilibrium with taxation than each one of them taken separately. For a given tax vector b , let \bar{f} be the social optimum with taxation and \hat{f} be the social optimum without taxation. There does not seem to be any a priori information about which of the two quantities

$$\frac{5}{4} \sum_e \bar{f}_e(l_e(\bar{f}_e) + b_e), \quad 2 \sum_e \hat{f}_e l_e(\hat{f}_e)$$

is smaller.

2 The Model

Let $G = (V, E)$ be a directed network (possibly with parallel edges but with no self-loops), and a set of *users*, each with an infinitesimal amount of traffic (flow) to be routed from an origin node to a destination node of G . Moreover, each user α has a positive *tax-sensitivity* factor $a(\alpha) > 0$. We will assume that the tax-sensitivity factors for all users come from a finite set of possible positive values. We can bunch together into a single *user class* all the users with the same origin-destination pair and with the same tax-sensitivity factor; let k be the number of different such classes. We denote by $d_i, \mathcal{P}_i, a(i)$ the total flow of class i , the flow paths that can be used by class i , and the tax-sensitivity of class i , for all $i = 1, \dots, k$ respectively. We will also use the term ‘commodity i ’ for class i . Set $\mathcal{P} \doteq \cup_{i=1, \dots, k} \mathcal{P}_i$. Each edge $e \in E$ is assigned a *latency function* $l_e(f_e)$ which

gives the latency experienced by any user that uses e due to congestion caused by the total flow f_e that passes through e . In other words, as in [5], we assume the additive model in which for any path $P \in \mathcal{P}$ the latency is $l_P(f) = \sum_{e \in P} l_e(f_e)$, where $f_e = \sum_{P \ni e} f_P$ and f_P is the flow through path P . If every edge is assigned a per-unit-of-flow tax $\beta_e \geq 0$, a selfish user in class i that uses a path $P \in \mathcal{P}_i$ experiences total cost of

$$\sum_{e \in P} l_e(f_e) + a(i) \sum_{e \in P} \beta_e$$

hence the name ‘tax-sensitivity’ for the $a(i)$ ’s: they quantify the importance each user assigns to the taxation of a path.

Let \hat{f} be a flow that satisfies the users’ demands and minimizes the social cost $\sum_{e \in E} f_e l_e(f_e) = \sum_i \sum_{P \in \mathcal{P}_i} f_P l_P(f)$, i.e., \hat{f} is a solution of the following mathematical program:

$$\begin{aligned} \min \quad & \sum_{e \in E} f_e l_e(f_e) \quad \text{s.t.} \quad (\text{MP}) \\ & \sum_{P \in \mathcal{P}_i} f_P = d_i \quad \forall i \\ & f_e = \sum_{P \in \mathcal{P}: e \in P} f_P \quad \forall e \in E \\ & f_P \geq 0 \quad \forall P \end{aligned}$$

Note that, although in certain cases (e.g. when the latency functions l_e are convex) the flow \hat{f} can be computed efficiently, for more general latency functions it may be extremely difficult to compute \hat{f} (see Section 4 in [5]). We will assume that such an \hat{f} exists and that it induces finite latency on every edge.

A function $g(x)$ is *positive* if $g(x) > 0$ when $x > 0$. We assume that the functions l_e are strictly increasing, i.e., $x > y \geq 0$ implies $l_e(x) > l_e(y)$, and that $l_e(0) \geq 0$. This implies that $l_e(f_e) > 0$ when $f_e > 0$, i.e., the function l_e is positive. Similar assumptions on monotonicity are made in [5].

Let

$$K := \{f : 0 \leq f_P, \forall P \wedge \sum_{P \in \mathcal{P}_i} f_P = d_i, \forall i\}$$

be the set of all flows that satisfy the users’ demands.

Definition 1. A flow f is called feasible iff $f \in K$.

A traffic (or Wardrop) equilibrium is a feasible flow $f^* \in K$ such that

$$\langle T(f^*), f - f^* \rangle \geq 0, \quad \forall f \in K. \quad (1)$$

where $\langle \cdot, \cdot \rangle$ denotes the inner product, and $T_P(f)$ is the function that gives the generalized cost of a user that uses path P when the network flow is f .